

Continuous Media on Demand



Researchers face the challenge of designing and implementing computer systems capable of processing and distributing large continuous media files simultaneously to millions of users who are connected to the Internet.

Jonathan C.L. Liu
University of Florida

David H.C. Du
University of Minnesota

Audio and video data—commonly referred to as “continuous media”—are characterized by their large data volume and continuous playback time, which typically requires storing them in a compressed form. The most popular digital audio and video compression formats include MPEG-1 (for video compact disks), MPEG-2 (for digital versatile disks), and MPEG-4 (for some wireless devices). Even compressed audio and video files can still represent a sizable data volume, and the challenge is to design and implement computer systems capable of storing, processing, distributing, and accessing these large continuous media files within heterogeneous communications environments.

With the rapid development of the Internet and the emergence of many new applications like e-learning, e-commerce, targeted advertising, digital television, and interactive television, many users will benefit from using continuous media. The accompanying sidebar by Richard H. Veith, “Interactive Video: Thirty Years and Counting,” provides a historical context for the shift from dedicated video-on-demand systems to Internet-based media-on-demand systems.

How to provide access to the desired continuous media simultaneously to the millions of users who are connected to the Internet at different places becomes a challenging issue. The key technology improvements, other than the Internet, that address this challenge include

- new compression methods that reduce data size yet maintain its quality,
- adequate client-side processing power to facilitate software decoding,

- the availability of broadband services to deliver high-quality continuous media data,
- advanced technologies for storing large volumes of data, and
- the widespread adoption of the client-server architecture to ensure the delivery of content across a broad geographic area.

This special issue explores the latest research on these aspects of continuous media on demand.

In “A Scalable and Reliable Paradigm for Media on Demand,” Gavin B. Horn and his colleagues discuss strategies for reliably transmitting unicast, multicast, or broadcast media over the Internet. They describe the use of forward error correction and a replication engine to reduce the number of packets transmitted over the network. “Streaming Technology in 3G Mobile Communication Systems” by Ingo Elsen and his coauthors gives an overview of third-generation mobile communication systems and the problems and challenges of streaming multimedia content. This is an area that is becoming increasingly important as content providers seek to meet the needs of users in the rapidly growing mobile market.

The performance of a continuous media server’s storage system typically bounds its scalability. Both the storage space capacity and the actual bandwidth limit the amount of data the server can store. Therefore, a single server can support only a limited number of accesses. To support a large number of concurrent accesses, it is clear that a number of servers must duplicate the data. If these servers are located in the same site, all accesses, regardless of their originated locations, must traverse a longer distance to reach the site. This tends to consume a large amount of back-

bone bandwidth, potentially creating network congestion problems. One alternative is to conserve network bandwidth by duplicating data in a proxy server—a server located in close proximity to the clients. The content a proxy server stores can be dynamically changed to reflect a local community’s needs. In “Proxy Servers for Scalable Interactive Video Support,” Husni Fahmi and colleagues discuss the idea of caching hotspots at proxy servers to reduce the load on video servers and decrease the random-seek response time. They also describe a simulation study in which they evaluated the hotspot-caching technique.

“Keyframe-Based User Interfaces for Digital Video” by Andreas Girgensohn, John Boreczky, and Lynn

Wilcox investigates the problem of how to find relevant information—for example, a specific video segment where a certain event occurred—in a video stream. The authors describe three visual interfaces to help users identify potentially useful or relevant video segments using keyframes—still images automatically extracted from video footage—and discuss corresponding user studies evaluating different design alternatives. In “Streaming-Media Knowledge Discovery,” Jan Pieper, Savitha Srinivasan, and Byron Dom point out the lack of a comprehensive resource list for locating streaming media information on the Internet and describe their work on building effective indexing and classification tools for locating streaming media data on the Internet.

Interactive Video: Thirty Years and Counting

Richard H. Veith, R. Veith Consulting

About 30 years ago, we first began to see technically feasible implementations of video on demand and two-way video, especially in the cable and closed-circuit TV environments, and in several video-phone attempts. The techniques used then included intermediate switching centers, local video storage, longer-than-real-time transmission, reduced bandwidth, transmitting only pixels that change, and various combinations of these methods. A 1976 survey provided the following examples of interactive video and video on demand from the early 1970s.¹

Rediffusion and GTE

In the mid to late 1960s, Rediffusion International of London developed a two-way video system that used intermediate switching exchanges. A wire pair between the home and the switching exchange had a 15-MHz bandwidth, and a second wire pair contained control signals. Subscribers used a rotary dial in the home to select among the channels available at the switching center. In a 1971 trial deployment in Dennisport, Mass., Rediffusion demonstrated two-way TV and a version of video on demand, with one subscriber sending video out on a channel that other subscribers would select by dialing.

Within the next year or so, in institutional installations in Europe, the US, and elsewhere, Rediffusion built systems that consisted of videocassette players arrayed in banks. Signal pulses from dialers in the rooms connected to the systems would select and start video transmissions from the players. In some of the installations, live video could be transmitted from any of the rooms, and any room could dial up a channel that contained a desired video transmission.

GTE developed a similar but more capable system, also in the early 1970s, with a 40-MHz coaxial cable between the home and a switching center. This system accommodated full-bandwidth TV transmissions in both directions simultaneously, which Rediffusion’s system could not do.

MITRE

At about the same time, the MITRE Corporation ran several different tests of interactive video using local frame stores. In a 1972 version for community colleges, the configuration included digital video memories in the student terminals that could store individual video frames and display them, as well as a provision for student terminals to select video from among 20 videocassette players at a central location.

A year or so earlier, MITRE used a cable TV system to test a similar setup for home viewers in Reston, Va. The system, based essentially on the transmission of individual frames stored on a videotape recorder in the home, could selectively address individual frames to specific home terminals. The system embedded the address in the video signal for each frame, in much the same way the BBC had begun to do with teletext.

Videophone

Videophone systems of the 1960s and 1970s used narrower bandwidths and functioned identically to what developers then called slow-scan television: The time required to transmit a picture frame would be longer than the time required to display it. Thus, picture information would accumulate in a local frame store until the system cached the entire frame.

Bell Labs began tinkering with two-way video in

We thank everyone who submitted papers, assisted in the review process, and helped to get the accepted papers into their published format. We hope you enjoy this special issue on media on demand. *

Jonathan C.L. Liu is an assistant professor in the Department of Computer and Information Science and Engineering at the University of Florida. His research interests include high-speed networks, multimedia communications, parallel processing, and artificial intelligence. Liu received a PhD in computer science and engineering from the University of Min-

nesota. He is a member of the ACM and the IEEE. Contact him at jcliu@cise.ufl.edu.

David H.C. Du is a US West Chair Professor in the Department of Computer Science and Engineering at the University of Minnesota and a founding member of Streaming21, a developer of video-streaming technologies enabling broadcast-quality video and audio on a large scale over the Internet. His research interests include high-speed networks, multimedia applications, high-performance computing over workstation clusters, database design, and CAD for VLSI circuits. Du received a PhD in computer science from the University of Washington. He is a member of the ACM and an IEEE Fellow. Contact him at du@cs.umn.edu.

1930, then demonstrated a rudimentary Picturephone system at the New York World's Fair in 1960. Commercial Picturephone service began in 1970.

In the early 1970s, RCA Global, Stromberg-Carlson, L.M. Ericsson, and others began offering alternative videophone systems or services. These systems used much narrower bandwidths than television, with an accompanying reduction in picture quality.

The time needed to accumulate a single picture frame varied. An early 1970s system by Robot Research of San Diego could transmit a picture every 8 seconds over normal telephone lines. The Bell System's Picturephone used three wire pairs and a video bandwidth of 1 MHz. The system digitally encoded the video transmission with a 3-bit code to convey only the part of the picture that changed, allowing a full frame to accumulate in only 3 seconds.

Warner Communications

A decade later, with personal computers and videodisk players on the market, the technology had changed somewhat, but the concepts remained essentially the same. In the early 1980s, long before Warner's Orlando trials—in which the company sought to link 4,000 subscribers to a full-service cable system that offered downloadable movies—a small group within Warner Communications explored the possibility of providing video on demand over cable TV systems. The concept included assembling banks of videodisk players at a central cable location, then using time-division techniques to route short clips of video to individual homes. During the intervals when a given home terminal was not receiving the real-time

video, a processor in the home terminal's computer would generate a simple computer graphics display showing a map or diagram that might lead to the selection of another video segment.²

Fast Forward

Today, we have more bandwidth, better compression techniques, immensely more computing power, the Internet, and all the advantages of a digital environment. In one sense, the past 30 years have seen video technology chasing the dream of successfully implementing the types of video on demand and two-way video services outlined in the late 1960s and early 1970s.

At the same time, as technology approached these goals, the playing field shifted. Now we want real-time HDTV over the Internet, wired and wireless, with video displays so real it appears as if we could touch the content. These desires have moved us from traditional television into the realm of complex computer systems and networks.

Although we appear to have come much closer to economically and technically feasible systems for residential video on demand and two-way video, we may have moved away from the idea that the average home user can access these services as easily and cheaply as the telephone and television.

References

1. R.H. Veith, *Talk-Back TV: Two-Way Cable Television*, Tab Books, Blue Ridge Summit, Pa., 1976.
2. R.H. Veith, *Visual Information Systems: The Power of Graphics and Video*, G.K. Hall, Boston, 1988, p. 151.